

SAM (REX/Ops) Operations Meeting 2008.08.12

Adam (typist), Robert, Eric, Art, Krzysztof, Angela, Anil, Steve

Meeting time: 35

Posting these minutes closes DZSAM-147

For reports submitted to the meeting...

The report text is in this font

Comments made during the meeting are in this font inline with the report

Releases:

DB server 8_5_1 (Robert's inaugural DB server) and client 8_5_0. Not deployed in production yet. One test production server at D0.

Changes in DB server are DB performance improvements stemming from DB problems two weeks ago.

Client has new DB server for uncaching files + minor fixes like Joel's modification of parameters (someone should tell Joel once it is installed).

Margaret complained about number of products in upd. There are 76 different version of the SAM station. Adam will remove most or all of these. But Art points out that UPD is an important archive of these executables. Why would anyone want to run "upd -aK+" anyway? It's all on the kits website.

DØ:

Notes (from shifters etc.):

--- Aug. 5 : downtime for d0ofprd1 and d0oflump (PRD) D0 offline databases which was completed before 11 am.

-- Aug. 9: A shifter reported that mezsilo library/994007 mover was in "yellow" for an entire (European shift).

-- Aug. 8: A shifter reported that samalto&samalto2 libraries/D21HLTO mover was in "yellow" for an entire (European shift).

-- Aug. 6: grid1_prd and grid2_prd DB servers restarted (by a shifter).

Were v5 ones that were busy thinking. Now there are only two v5 DB servers. No one has complained.

Unresolved issues (of the past week):

-- DZSAM 208: I almost wanted to close this ticket. The Listserv complained that "d0sam-onshift" has sent too many email

and our posting privilege has been withdrawn for the next 48 hours ? ... Wonder anything needs to be done or not.
Not sure what this is. Not sure who is sending all of the mail.

-- DZSAM 207: "samgrid.fnal.gov:8080/list_of_schedulers.php disappeared and then reappeared ? Do we know why ?

Was probably the periodic expression period

-- DZSAM 205: Problem with samgrid web server or not ?

-- DZSAM 203: J. Snow complained about "OSG MC jobs being held".

-- DZSAM 201: "MC merge jobs fail at MWT2-IU" --- Should this ticket be closed (as the problem has seemed to disappear) ?

This was where Joel complained about weird CORBA errors, but the errors have stopped. We think things are working again.

Add note about samgrid reaper...

CAB cache nodes are being moved from FCC to GCC. Ask for the plan. Since most people are in Prague, the reduction in service should not be noticed.

CDF:

I've been on CDF SAM shift during last week, here's my report.

The shift has been really quiet over all week, here are some minor issues encountered:

fcdfdata147 offline, helpdesk ticket pending.

---Fri Aug 8

Tape IAB639 marked no access and LTO3_22 Mover offline:

Comment by Stanley Hicks:

=====

The volume was marked NOACCESS and the drive went offline when it was too long in SEEK and enstore tried to remount it in different drives without first unmounting it in the current drive.

There is a bug in the program which has been fixed and will be installed on all movers with the next downtime for cdf (supposedly next Wed).

=====

---Sat Aug 9

Many Enstore alarms like this first message:

=====

Fri Aug 8 23:41:00 CDT 2008

From: alarm_server

['1218256860.93', 1218256860.927525, 'fcdfdata106.fnal.gov',
27540, 'root', 'C (1)', 'ENCP',
'DEVICE ERROR', None, None, None, {'text': {'status': '[Error 16]
Device or resource busy: write()
system call stuck inside kernel too long', 'outfile':
'/export/data2/read-pool-2/data/001700000000000006B0A88', 'infile':
'/pnfs/fs/usr/.(access)(001700000000000006B0A88)'}, 'r_a':
(('131.225.211.98', 51042), 1L,
'131.225.211.98-51042-1218255979.192790-27540')}]

=====

Tape IAB762 marked NO ACCESS.

Comment by Stanley Hicks:

=====

This volume was too long in seek and then went NOACCESS because enstore
tried mounting it in different drives to read it. It was marked full on

2007-06-28 02:25:59.404751 system_inhibit[1] full

Some new CDF dCache nodes were installed with SL5 but did not work. Were
installed with SL4 instead (Adam dig out CD Ops report and add here)...

MINOS:

Scaling up PARROT tests for analysis on FermiGrid,
using shared local cache, plus squid.fnal.gov .

Parrot allows access of as from any grid node (via http). Art will test. Need to unravel
symbolic links to make Parrot happier.

Enstore -

DCache -

SAM -

Oracle -

OPEN ISSUES

IT 1146: cannot test existence of sam tape location

IT 2071: quadratic project startup time growth

PLAN -

DB:

D0 :

- DB security patches on d0ofint1 and d0ofprd1 for Aug 05, 2008 - Done

Alert(S):

CDF :

- Scheduling DB security patches on cdfofdev/cdfofi nt on Tuesday 19th Aug.

There will be small interruption to cdfofdev on Monday 18th in order to put the db in archive mode so that

backup will be taken before patch.

When to schedule the production db patch? Angela will find out.

Alert(s):

MINOS:

Alert(s):

Enstore Operations report (from CD Ops report)

- Release Management:

+ Scheduled Downtimes:

On Wednesday 8/13 at 8:30 AM we would like to upgrade the Movers and Library Manager code on cdfen This is about a 30 minute downtime and can be done with minimal interruption to normal processing.

- Problem Management:

+ Unscheduled Downtimes:

We had two occurrences of a cap failure on the FCC SL8500. A tape was not fully seated after the cap rotated to the inward position so the Bots could take tapes from the cap. We are not sure if these are a Bot

gripper issue or a bad cap or just an accident while loading the cap. We are tracking the issue. This resulted in about 3.5 hours total downtime on the specific robot.

- Capacity Management:

CDF rack of movers is here. We are installing and testing. Should be able to turn them over later this week for CDF non-RAW and GP stken.

Tape drive and mover requisitions are in Procurement, checking on progress. Bids for movers are due in tomorrow.

SL8500 requisitions are in Procurement, checking on progress.

Loaded 600 CDF LTO4 tapes into SL8500#3.

- Incident Management:

Data Integrity:

Historical report (a work in progress):

7/14	7/21	7/28	8/4	8/11	
17	21	48	5	25	enstore tape mover interventions
20	34	50	30	24	tape incidents
1	1	1	0	2	library hardware incidents
0	0	1	1	0	tape drive hardware incidents
0	0	0	5	1	enstore mover hardware incidents
0	0	2	1	0	enstore server hardware incidents
2	0	0	0	0	dcache pool hardware incidents
0	0	0	0	0	dcache server hardware incidents
0	0	0	0	0	pnfs hardware incidents
3	3	0	3	0	dcache software incidents
0	1	0	2	0	enstore software incidents
0	0	0	0	0	enstore configuration incidents
1	0	0	0	0	dCache configuration incidents
7	4	0	0	3.5	unsched. enstore outage hours
0	0	0	0	0	unsched. dCache outage hours

Metrics of Enstore planned and unplanned (partial) downtimes for all of Enstore:

[http://www-ccf.fnal.gov/IA_metrics/Running_Sum_Planned_Downtime\(hours\).jpg](http://www-ccf.fnal.gov/IA_metrics/Running_Sum_Planned_Downtime(hours).jpg)
[http://www-ccf.fnal.gov/IA_metrics/Running_Sum_All_Downtime\(hours\).jpg](http://www-ccf.fnal.gov/IA_metrics/Running_Sum_All_Downtime(hours).jpg)

- Migration work flow status:

We have 4 migration servers in use now. 3 are migrating tapes, 1 is scanning.

LTO1 Migration: 0 / ~1200
LTO2 Migration: 1676 / ~8500 (many duped but no final scan yet)
9940A Migration: 1726 / ~4500
9940B Migration: 0 / ~24000

D0 migration progress plots:

http://www-d0en.fnal.gov/enstore/migration_summary/migration_summary_daily_3480.jpg
http://www-d0en.fnal.gov/enstore/migration_summary/migration_summary_accumulated_3480.jpg
http://www-d0en.fnal.gov/enstore/migration_summary/migration_summary_daily_9940.jpg
http://www-d0en.fnal.gov/enstore/migration_summary/migration_summary_accumulated_9940.jpg

- SSA group is interviewing for a Full time position. Interviews and phone screens in progress.

FNDCA

- no dcache issues reported

CDFDCA

- cdf dh sent requested regarding old log files gzipped by Timur to release space on cdfdcache5:/export . I removed few files but situation was taken care already. In general, the disk space on root area on cdf dcache admin nodes is low (total disk space is ~40GB). Is suggested cdf dh to address disk space issue on these nodes. They will be replacing admin nodes in the the fall.

- Alex has configured and put several new dCache pools in production. pool nodes upgrade:

- fcdldata320 has dcache reinstalled and ready to go waiting for relocation to GCC.
- fcdfcache1-8 +9 : installed dcache, there are issues with upd/ups encp installation. run2sys reported some files were not installed properly (SLF5), and run2sys will reinstall OS on these nodes. Requested to roll back OS to SLF 4.4 - we do not have dcache running on SLF5.x here at FNAL, pnfs unmounts cleanly but only with "-force" option. Few other things needed to be taken care of by run2sys and/or CDF DH.
- at present a few restore requests are failing go to tape with noaccess error.
- dCache reinstalled on pool 098 to restore pool disk configuration after disk repair. Node back in production.